

Appendix DRemote Network Communication FacilitiesLimitations for Interactive Work

Users asked to accept a remote computer as if it were next door will use a local telephone call to the computer as a standard of comparison. Current network terminal facilities do not fully accomplish the illusion of a local call. Data loss is not a problem in most network communications - in fact with the more extensive error checking schemes, data integrity is higher than for a long distance phone link. On the other hand, networking relies upon shared community use of telephone lines to procure widespread geographical coverage at substantially reduced cost. Unless enough total line capacity is provided to meet peak loads, substantial queueing and traffic jams result in the loss of terminal responsiveness. Limited responsiveness for character-oriented TENEX interactions continues to be a special problem for network users and is one of the reasons that coming more local computing systems will be especially important to improve the human interfaces to our AI programs. The key technological components to improved human engineering (high-speed bit-mapped displays, touch, and speech) all involve requirements for high bandwidth communications that can only be effectively implemented locally.

This does not diminish the importance of networks in our community, but rather enhances their role for facilitating remote scientific contacts, allowing remote access to regionalized resources, and sharing programs and knowledge bases. These are tasks for which national networks are ideally suited.

TYMNET

TYMNET provides broad geographic coverage for terminal access to SUMEX, spanning the country and also increasingly accessible from foreign countries (see Figure 18 on page 379). TYMNET has made few technical changes to their network that affect us other than to broaden geographical coverage. The previous network delay problems are still apparent although better cross-country trunks into New York and New England are installed and improving service there. TYMNET is still primarily a terminal network designed to route users to an appropriate host and more general services such as outbound connections originated from a host or interhost connections are only done on an experimental basis. This presumably reflects the lack of current economic justification for these services among the predominantly commercial users of the network. Whereas TYMNET is developing interfaces meeting X.25 protocol standards, the internal workings of the network will likely remain the same, namely, constructing fixed logical circuits for the duration of a connection and multiplexing characters in packets over each link between network nodes from any users sharing that link as part of their logical circuit.

We have continued to purchase TYMNET services through the NLM contract with TYMNET, Inc. Because of current tariff provisions, there is no longer an economic advantage to this based on usage volume. SUMEX charges are computed on its usage volume alone and not the aggregate volume with NLM's contribution to achieve a lower rate. A new tariff provision, based on "dedicated port" pricing, is advantageous to us though. This allows purchase of a number of logical network ports at the host for a fixed cost per month, independent of connect time or number of characters transmitted. We have implemented that option with BRP and save approximately \$1,000 per month in service charges. We will continue to work closely with NIH-BRP and NLM to achieve the most cost-effective purchase of these services. The total use of TYMNET dropped during the TELENET experimental connection described below (see Figure 16) but has increased again since the TELENET service was dropped.

Technical aspects of our connection to TYMNET have remained unchanged since the last report and have continued to operate reasonably reliably. We have fixed several bugs in the TYMNET service related to handling editing terminals. Also we have had problems with incomplete closure of connections that can accumulate and leave us with all ports effectively blocked after long periods of uptime. The evidence points to a bug in TYMNET's interface code and we have had serious problems getting adequate support from them to fix the problem.

ARPANET

We continue our advantageous connection to the Department of Defense's ARPANET, now managed by the Defense Communications Agency (DCA). Current ARPANET geographical and logical maps are shown in Figure 19 and Figure 20 on page 380. Consistent with agreements with ARPA and DCA we are enforcing a policy that restricts the use of ARPANET to users who have affiliations with DoD-supported contractors and system/software interchange with cooperating network sites. We have maintained good working relationships with other sites on the ARPANET for system backup and software interchange. Such day-to-day working interactions with remote facilities would not be possible without the integrated file transfer, communication, and terminal handling capabilities unique to the ARPANET. The ARPANET is also key to maintaining on-going intellectual contacts between SUMEX projects such as the Stanford Heuristic Programming Project authorized to use the net and other active AI research groups in the ARPANET community.

The reconnection of the Rutgers resource to ARPANET has reopened our valuable scientific contacts with that subcommunity. In fact their efforts to justify reconnection may provide a basis for broader NIH use of the ARPANET and hence better network support for our collaborators.

TELENET

Initially SUMEX based its remote communication services on two networks - TYMNET and ARPANET. These were the only networks existing at the start of the project which allowed foreign host access. A third commercial network system, TELENET, is now competitively operational and offers a growing selection of services. Since our last review and with the advice and approval of the AIM Executive Committee and NIH-BRP, we established an experimental connection to TELENET to evaluate its technical and economic advantages relative of our existing connections. This initial experiment was unsuccessful but since then TELENET has been acquired by General Telephone and Electronics to provide a larger capital base. They have an aggressive program for augmenting network services and a reconnection may be of advantage sometime in the next grant term. A current TELENET network map is shown in Figure 21 on page 382.

Our experimental connection was via a TP-2200 interface with 12 asynchronous lines to the SUMEX host and one 4800 baud line connecting to the network proper. TELENET has many attractive features in terms of a symmetry analogous to that of the ARPANET for terminal traffic and file transfers and being a commercial network, it does not have the access restrictions of the ARPANET. Its tariff schedule also affords lower costs than TYMNET for comparable service volume.

However, despite system changes we made to optimize TELENET performance (Xon/Xoff facilities to improve traffic flow), users felt a substantial degradation in service when using TELENET as opposed to TYMNET. We insisted that users use TELENET whenever possible between November 1978 and May 1979 to maximize user accommodation so that problems arising from differences in access conventions would not cloud judgements of services. Complaints included poor node reliability, intolerable delays in response, uneven flow of terminal output, and poor operational management of the network in keeping users informed of network and host status. From the system viewpoint at SUMEX, we detected similar problems. We received ineffective system engineering support in trying to tune network parameters to optimize performance for our user community and poor or erroneous feedback about network failures and problem resolution. In practice, TELENET offered no service advantages over TYMNET, since no file transfer connections above 1200 baud were allowed, no facilities to control local versus remote echoing existed, and no electronic mail system existed to facilitate communication between network operations staff and host nodes. Also company financial problems portended substantial delays in remedying these problems.

Because of grant budget limitations, we were forced to decide between the TYMNET and TELENET connections. Based on the distinct user preference expressed for TYMNET, we decided to terminate the TELENET connection as of May 1, 1979. We will continue to monitor TELENET developments (and those of other potential national network servers, e.g., AT&T, IBM, and Xerox) and may recommend a reevaluation of an alternative source for network services in the future.

Figure 18a. TYMNET Network Access

TYMNET®**Domestic Access
Locations**

State	High density locations	Low density locations	Foreign exchange locations
Alabama		Birmingham	
Arizona		Phoenix* Tucson*	
Arkansas			Little Rock
California	El Segundo* Los Angeles* Mountain View Newport Beach* Oakland* Palo Alto Riverside/Colton San Francisco* San Jose/Cupertino* Ventura/Oxnard	Hayward Sacramento San Diego* Santa Rosa	Alhambra Burlingame Fresno Marina del Rey Norwalk San Clemente San Pedro Santa Barbara Van Nuys
Colorado	Denver*		Colorado Springs
Connecticut		Darien* Hartford*	Bridgeport Danbury New Haven Waterbury
Delaware			Wilmington
District of Columbia	Washington*		
Florida		Jacksonville Miami* Orlando* St. Petersburg Tampa	Ft. Lauderdale Pensacola West Palm Beach
Georgia		Atlanta*	Savannah
Idaho		Boise	
Illinois	Chicago*	Freeport Rockford Springfield*	Peoria
Indiana		Indianapolis* South Bend	Evansville Ft. Wayne Marion

*1200-baud access

September 1979

State	High density locations	Low density locations	Foreign exchange locations
Iowa		Des Moines Iowa City	Cedar Rapids
Kansas		Shawnee Mission* Wichita*	Topeka
Kentucky		Lexington Louisville*	
Louisiana		Baton Rouge* New Orleans*	Lafayette Shreveport*
Maryland	Baltimore*		
Massachusetts	Boston/Cambridge*	Springfield	
Michigan	Ann Arbor Detroit* Plymouth Southfield	Jackson Kalamazoo	Grand Rapids St. Joseph
Minnesota		Minneapolis*	
Mississippi			Jackson
Missouri		Kansas City* St. Louis*	
Nebraska		Omaha	
Nevada		Reno/Carson City*	Las Vegas
New Hampshire			Manchester Nashua
New Jersey	Englewood Cliffs Lyndhurst* Newark/Union* Piscataway Wayne	Princeton/ South Brunswick	Moorestown
New Mexico			Albuquerque*
New York	New York City*	Buffalo* Corning Rochester* Syracuse White Plains*	Albany Hempstead L.I. Huntington L.I. Niagara Falls

Figure 18a. TYMNET Network Access

(continued)

State	High density locations	Low density locations	Foreign exchange locations
North Carolina		Raleigh/Durham Winston-Salem*	Charlotte Greensboro
Ohio		Akron Cincinnati Cleveland* Columbus Dayton	Toledo
Oklahoma		Oklahoma City* Tulsa*	
Oregon		Portland*	
Pennsylvania	Philadelphia*	Erie Pittsburgh* Valley Forge	Allentown Harrisburg York
Rhode Island			Providence
South Carolina			Columbia Greenville
Tennessee		Chattanooga Memphis* Nashville	Knoxville
Texas	Houston*	Austin* Dallas* El Paso* Midland San Antonio	Baytown Beaumont Corpus Christi Ft. Worth Longview Lubbock Odessa
Utah		Salt Lake City*	
Vermont			Burlington
Virginia	Arlington*		Norfolk Richmond
Washington		Richland* Seattle*	Spokane*
West Virginia			Charleston
Wisconsin		Madison Milwaukee*	Oshkosh

TYMNET®

International Access Locations

Argentina
Buenos Aires

Australia*
Sydney

Austria
Vienna

Bahrain

Belgium
Brussels

Bermuda†

Brazil
Rio de Janeiro

Canada
All Datapac cities

Denmark
Copenhagen

Finland
Helsinki

France*
Paris

Germany*
Frankfurt

Hong Kong

Israel
Tel Aviv

Italy
Milan
Rome

Japan†
Tokyo

Mexico
Mexico City

Netherlands
Amsterdam

New Zealand†
Wellington

Norway
Oslo

Philippines
Manila

Portugal
Lisbon

Puerto Rico
San Juan

Singapore
Singapore

Spain*
Madrid

Sweden
Farsta

Switzerland*
Berne

United Kingdom
London

United States†
Anchorage
Honolulu
Juneau

* Access can be made throughout the country with a local call.

† Projected for 1980.

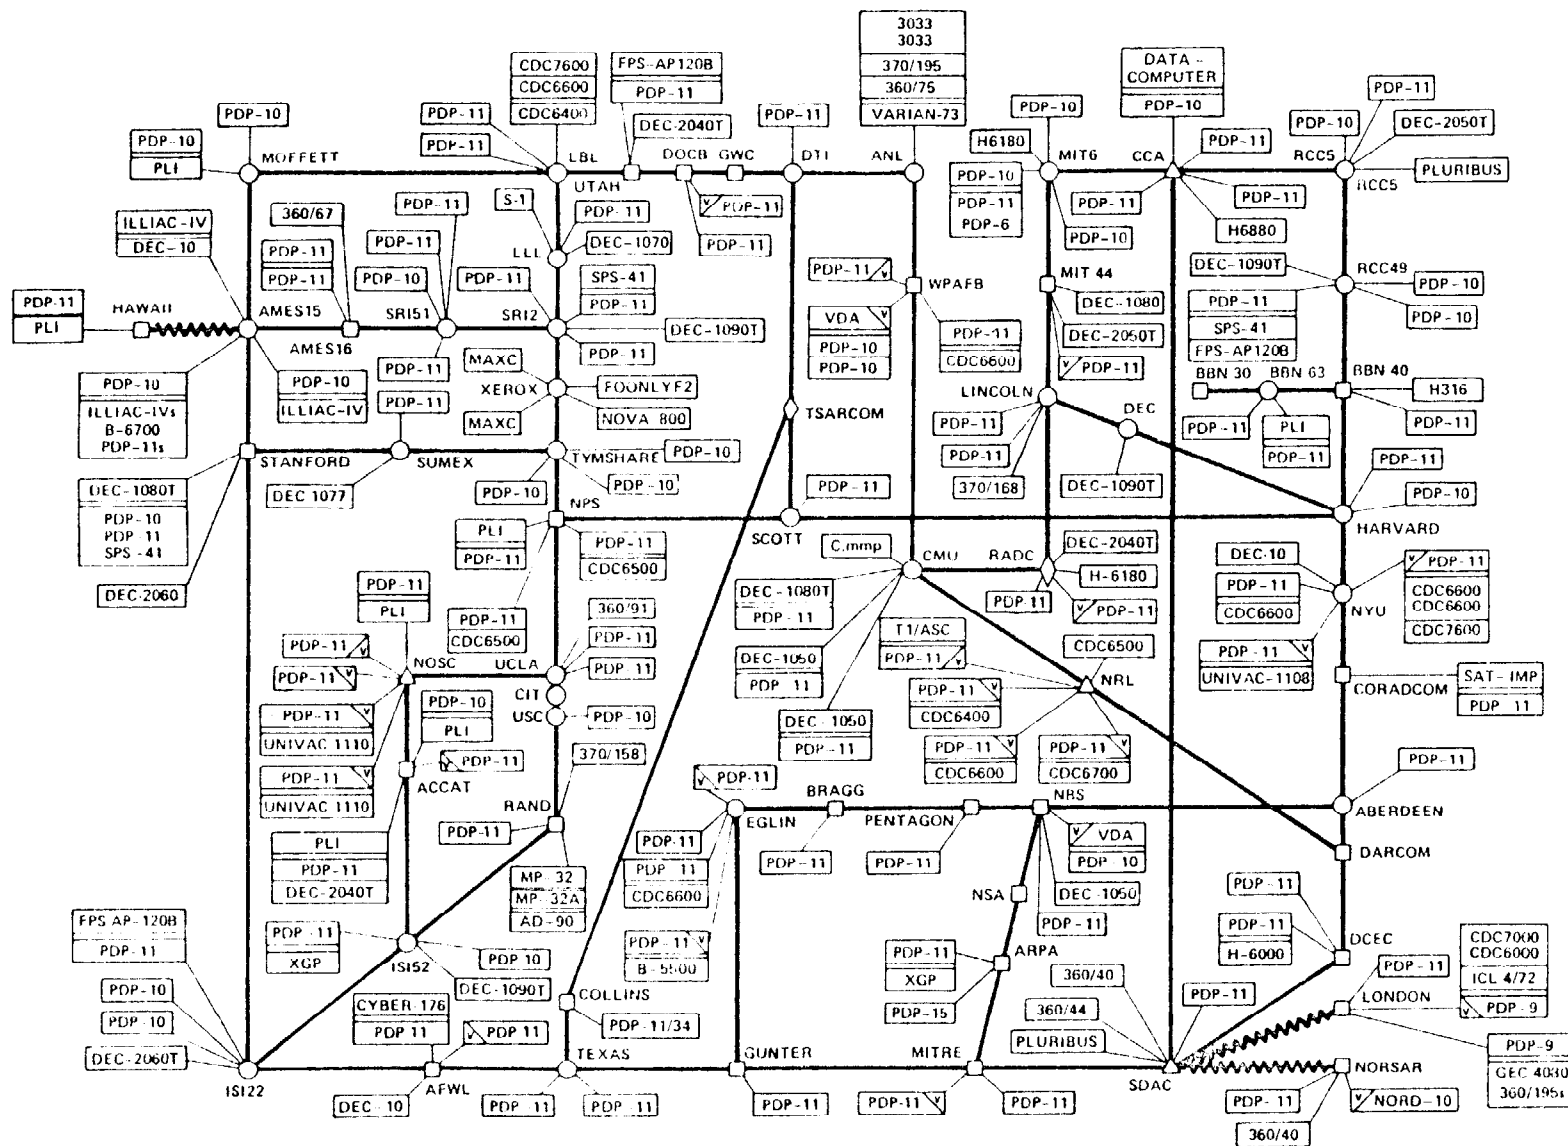
‡ Noncontinental.

October 1979

Figure 18b. TYMNET Network Access


ARPANET LOGICAL MAP, MARCH 1980

Figure 20. ARPANET Logical Network Map



- IMP
- TIP
- △ PLURIBUS IMP
- ◇ PLURIBUS TIP
- ~~~~~ SATELLITE CIRCUIT
- ▽ VERY DISTANT HOST

PLEASE NOTE THAT WHILE THIS MAP SHOWS THE HOST POPULATION OF THE NETWORK ACCORDING TO THE BEST INFORMATION OBTAINABLE, NO CLAIM CAN BE MADE FOR ITS ACCURACY
 HOST COMPUTER CONFIGURATION SUPPLIED BY THE NETWORK INFORMATION CENTER
 NAMES SHOWN ARE IMP NAMES, NOT (NECESSARILY) HOST NAMES



(CURRENT) THE TELENET NETWORK

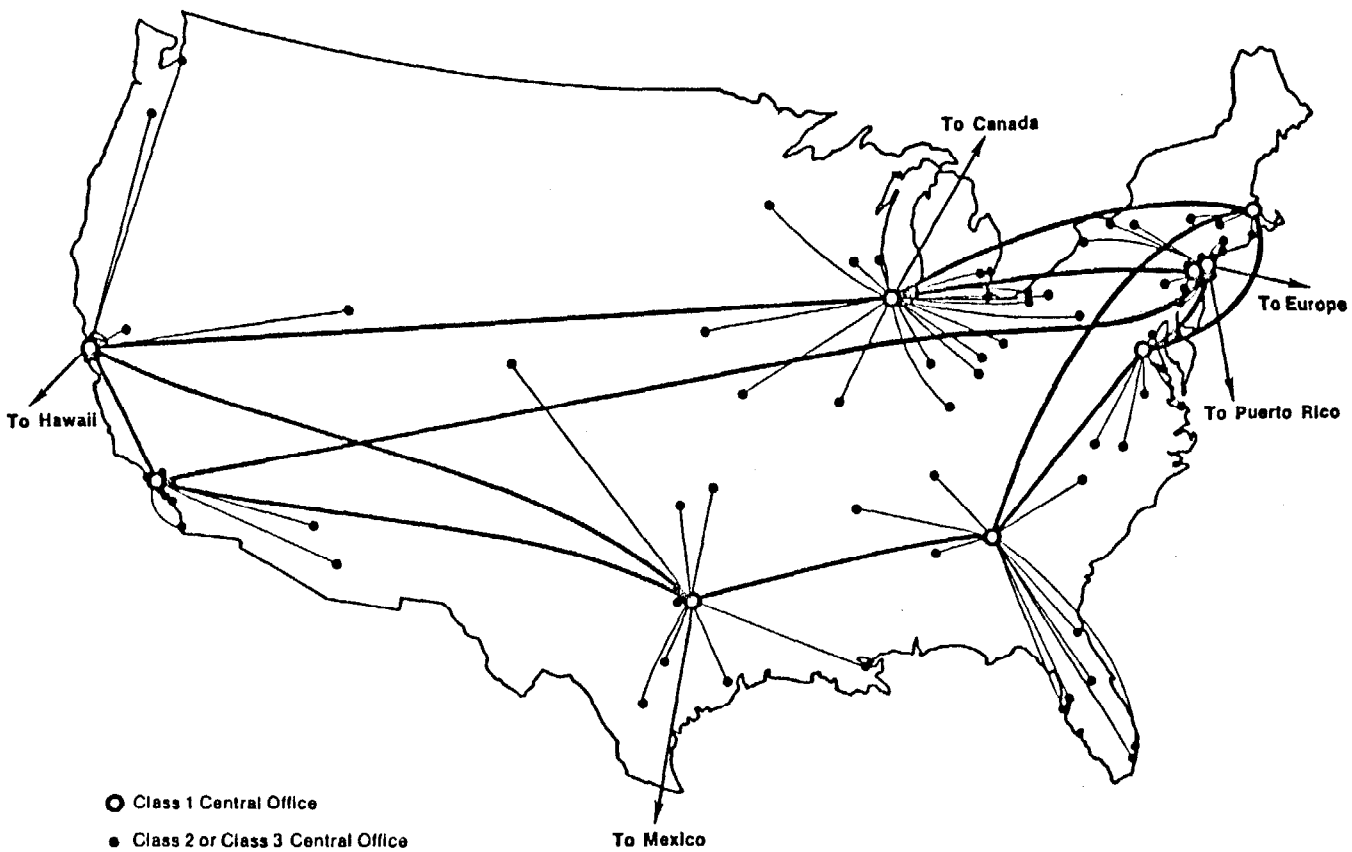




Figure 21a. TELENET Geographical Network Map

Figure 21b. (MID 1980)

THE TELENET NETWORK



Appendix EResource Management StructurePhilosophy of Management

One way to administer a national resource is by subcontract to a fee-compensated, neutral agent under a governing body that could speak to the technical and quality-control interests of the served constituency. Appropriate in some circumstances, this model would separate the administration of the resource from active participation in the on-going research and development. An approach expected to foster greater creativity is to couple the resource closely with an active user-center. This of course can lead to manifest conflicts of interest that must be addressed and avoided if the resource is to be available fairly on a regional or national basis.

SUMEX-AIM has been based on the latter approach with a charter that spells out the underlying objectives and responsibilities of the program, and which establishes incentives, resources, and obligations for proper performance. Our resource design, incorporating all of these ingredients, has made the development of the procedural framework a matter of simple common-sense logic. It will be plain that the convergence of local self-interest with peer and contractual responsibility offers the best assurance that the programmatic goals will be respected and simplifies the tasks of surveillance and accountability.

The self-interest part of this equation stems from our original motivation in requesting the resource: the need for specialized computing facilities to support intense, interdisciplinary studies in applications of AI at Stanford University Medical School. Comprising several departments (Chemistry, Medicine, Genetics, and Computer Science), interwoven projects (DENDRAL, MYCIN, MOLGEN, Heuristic Programming), and principal faculty (Professors Feigenbaum, Lederberg, Djerassi, Shortliffe, and Buchanan), a substantial body of research has progressed and evolved over many years. Successful, stable collaborations of this scope are not readily found. This history both depends upon and contributes to the doctrine of resource-sharing that underlies the SUMEX-AIM effort.

One premise of the management plan is therefore the charter allocation of half the user-available capacity of the SUMEX facility to the Stanford complex of projects, subject to a local committee chaired by Professor Feigenbaum. This principle clearly defines the local benefit of the resource, minimizes anxiety and conflict-of-interest, and enables the local group to respond quite objectively to the allocations that are made by an Executive Committee for the "national" or non-Stanford aliquot (see the section on "Management Committees" below). Another important contribution to the success of the plan is the welcome participation of an NIH-BRP representative on the Executive Committee. What would be inappropriate meddling in the conduct of a narrower research project funded by NIH, is a communication channel and source of detached judgment that has

been invaluable in expediting the innumerable decisions about which NIH must and should be consulted in the week-to-week business of the resource. The efficacy of this principle, as is appropriate to acknowledge here, has been validated and enhanced by the style and energy that Dr. William Baker has brought to this task.

Further consequences of the charter principles are the conscientious cultivation of the "national" community for the most efficacious use of its aliquot, and the further growth of distributed facilities in due course. In summer of 1977, a computing facility at Rutgers University was established, coupled to SUMEX-AIM via the ARPANET and with 15% of the user-available capacity allocated for AIM use with the advice of the AIM Executive Committee. An increasing number of projects are using that resource as reported in Section 9.

Finally, the recognition in the charter that SUMEX-AIM is not merely a retail-store for computer cycles, but the means of building a community, is a necessary basis for the morale of the whole operation and the rationale for no fee-for-service.

The remainder of this section will summarize the way in which these responsibilities are handled bureaucratically.

Organization and Procedures

The SUMEX-AIM resource is administered between the Departments of Medicine and Computer Science of Stanford University. Its mission, locally and nationally, entails both the recruitment of appropriate research projects interested in medical AI applications and the catalysis of interactions among these groups and the broader medical community. User projects are separately funded and autonomous in their management. They are selected for access to SUMEX on the basis of their scientific and medical merits as well as their commitment to the community goals of SUMEX. Currently active projects span a broad range of application areas such as clinical diagnostic consultation, molecular biochemistry, psychological and affective behavior modeling, instrument data interpretation, and tool building to facilitate the development of new AI applications.

In July 1978, Professor Lederberg, the original SUMEX Principal Investigator, became president of The Rockefeller University. Professor Feigenbaum, chairman of the Stanford Department of Computer Science, took over as Principal Investigator of the SUMEX project. Because of Prof. Feigenbaum's role as co-Principal Investigator of SUMEX from its start and his long standing collaboration with Prof. Lederberg, the management transition took place very smoothly. The SUMEX-AIM community continues to function with the same high level of vitality as before and has continued to grow. Professor Lederberg retains an active role in the SUMEX-AIM community as chairman of the AIM Executive Committee and on a more frequent basis through the system message facilities.

Close scientific and administrative ties are retained with the Stanford medical community. Immediately following Prof. Lederberg's

departure, Professor Stanley Cohen, new chairman of the Department of Genetics, provided this liaison. In recognition of the growing scope and significance of the clinical applications being pursued at SUMEX, we have recently significantly strengthened our contacts within the Stanford community in that area. Professor Edward H. Shortliffe, one of the key designers of MYCIN, has assumed the role of co-Principal Investigator of SUMEX and the project will become administratively part of the Stanford Department of Medicine, effective August 1980. As part of the largest clinical medicine department at Stanford, SUMEX will have increased visibility and opportunity to broaden its local scientific collaborations.

Management Committees

Since the SUMEX-AIM project is a multilateral undertaking by its very nature, we have created several management committees to assist in administering the various portions of the SUMEX resource. As defined in the SUMEX-AIM management plan adopted at the time the initial resource grant was awarded, the available facility capacity is allocated 40% to Stanford Medical School projects, 40% to national projects, and 20% to common system development and related functions. Within the Stanford aliquot, Prof. Feigenbaum has established an advisory committee to assist in selecting and allocating resources among projects appropriate to the SUMEX mission. The current membership of this committee is listed in Appendix I.

For the national community, two committees serve complementary functions. An Executive Committee oversees the operations of the AIM resources (SUMEX and the AIM portion of the Rutgers facility) as related to national users and makes the final decisions on authorizing admission for new projects and revalidating continued access for existing projects. It also establishes policies for resource allocation and approves plans for resource development and augmentation within the national portion of SUMEX (e.g., hardware upgrades, significant new development projects, etc.). The Executive Committee oversees the planning and implementation of the AIM Workshop series implemented under Prof. S. Amarel of Rutgers University and assures coordination with other AIM activities as well. The committee will play a key role in assessing the possible need for additional future AIM community computing resources and in deciding the optimal placement and management of such facilities. The current membership of the Executive committee is listed in Appendix I.

Reporting to the Executive Committee, an Advisory Group represents the interests of medical and computer science research relevant to AIM goals. The Advisory Group serves several functions in advising the Executive Committee; 1) recruiting appropriate medical/computer science projects, 2) reviewing and recommending priorities for allocation of resource capacity to specific projects based on scientific quality and medical relevance, and 3) recommending policies and development goals for the resource. The current Advisory Group membership is given in Appendix I.

These committees have functioned actively in support of the resource. Except for the meetings held during the AIM workshops, the committees have "met" by messages, net-mail, and telephone conference owing to the size of the groups and to save the time and expense of personal travel to meet face to face. The telephone meetings, in conjunction with terminal access to related text materials, have served quite well in accomplishing the agenda business and facilitate greatly the arrangement of meetings. Other solicitations of advice requiring review of sizable written proposals are done by mail.

New Project Recruiting

The SUMEX-AIM resource has been announced through a variety of media as well as by correspondence, contacts of NIH-BRP with a variety of prospective grantees who use computers, and contacts by our own staff and committee members. The number of formal projects that have been admitted to SUMEX has nearly quadrupled since the start of the project; others are working tentatively as pilot projects or are under review. Reports for the various projects can be found in Section 9 and a graphical summary of community growth in Appendix B.

In the recent past we have made numerous efforts to broaden outside awareness of work in the AIM community and to encourage new research projects including:

- 1) CONGEN workshop at Stanford, December 1978.
- 2) AGE workshop at Stanford, February 1980.
- 3) AI session in the Fourth Illinois Conference on Medical Information Systems, 1979.
- 4) INTERNIST participation in a course on AI computing at NIH, 1979.
- 5) AI session in the Association for Information Science meeting, 1979.
- 6) AI session at Sixth International Joint Conference on AI, August 1979 and extensive lecture tour among Japanese university and industrial research projects.
- 7) MYCIN and INTERNIST program demonstrations at the American College of Physicians meetings in 1979 and 1980.

We have prepared a variety of materials for prospective new users ranging from general information in a SUMEX-AIM overview brochure to more detailed information and guidelines for determining whether a user project is appropriate for the SUMEX-AIM resource. Dr. E. Levinthal has prepared a questionnaire to assist users seriously considering applying for access to SUMEX-AIM. Pilot project categories have been established both within the Stanford and national aliquots of the facility capacity to assist and encourage new projects in formulating possible AIM proposals and pending their application for funding support. Pilot projects are approved for

access for limited periods of time after preliminary review by the Stanford or AIM Advisory Group as appropriate to the origin of the project.

These contacts have sometimes done much more than support already formulated programs and have provided guidance for new investigators and projects to formulate new biomedical AI applications and establish appropriate collaborations between medical and AI scientists. The AIM Executive and Advisory Committees have also played important roles in suggesting to pilot efforts ways in which their research programs could be strengthened through better collaborative ties.

We have welcomed a number of visiting investigators at Stanford who were able to pay their own expenses, so they could see first hand how AI applications programs are formulated and get acquainted with the computing tools available. As an additional aid to new projects or collaborators with existing projects, we provide a limited amount of funds for use to support terminals and communications needs of users without access to such equipment.

Stanford Community Building

The Stanford community has undertaken several internal efforts to encourage interactions and sharing between the projects centered here. Numerous classes and seminars have been held over the years including ones to introduce chemistry students to the DENDRAL programs and to develop the early versions of the AI Handbook 5 articles. We also hold weekly informal lunch meetings (SIGLunch) between community members to discuss general AI topics, concerns and progress of individual projects, or system problems as appropriate as well as having frequent outside invited speakers.

Existing Project Reviews

We have conducted a continuing careful review of on-going SUMEX-AIM projects to maintain a high scientific quality and relevance to our biomedical AI goals and to maximize the resources available for newly developing applications projects. At the last full AIM workshop, meetings of the AIM Advisory Group and Executive Committee were held to review the national AIM projects. These groups recommended continued access for all formal projects then on the system. They also recommended phasing out the Organ Culture pilot project.

In the fall of 1978, meetings of the Stanford Advisory Group were held to review projects supported out of the Stanford aliquot. The recommendation of this group was to phase out support for the Hydroid Project, pending work more directly applicable to SUMEX-AIM goals. The group also recommended phasing out the Quantum Chemistry and Genetics Applications pilot projects unless stronger AI relevance were established immediately. The Quantum Chemistry project has since developed close collaboration with the DENDRAL stereochemistry effort. The Genetics Applications project has transferred their work to other systems to continue their calculations on genetic demographic data and has stopped using SUMEX.

AIM Workshop Support

The Rutgers Computers in Biomedicine resource (under Dr. Saul Amarel) has organized a series of workshops devoted to a range of topics related to artificial intelligence research, medical needs, and resource sharing policies within NIH. Until recently, meetings have been held regularly at Rutgers.

In May 1979, a mini-AIM workshop devoted to clinical diagnosis programs was organized by MIT-Tufts and Rutgers and held in Vermont. This meeting was small (about 25 attendees) and emphasized detailed technical discussions about system designs and the strengths and weaknesses of various approaches. Many of the attendees were graduate students in order to maximize the benefit of personal contacts and discussions for on-going research projects. Topics covered in the discussions included state-of-the-art in explanation, causality in reasoning, strategies of focusing and dealing with multiple diagnostic problems, issues of representation and grain of description, creating and updating a knowledge base, planning strategies, issues of time representation, and inexact reasoning.

In August 1980, the AIM workshop will be held at Stanford as part of an extensive series of meetings. The workshop will be followed by a two-day series of tutorials for medical scientists to introduce them to AI computing goals and capabilities. This in turn will be followed by the first annual conference of the American Association for Artificial Intelligence devoted to a broad range of scientific issues in AI research.

The SUMEX facility has served as a communications base for workshop planning and provided support for workshop demonstrations when requested. We expect to continue this support for future workshops. The AIM workshops provide much useful information about the strengths and weaknesses of the performance programs both in terms of criticisms from other AI projects and in terms of the needs of practicing medical people. We plan to continue to use this experience to guide the community building aspects of SUMEX-AIM.

Resource Capacity Planning and Allocation Policies

As the SUMEX-AIM community has grown, the facility has become increasingly loaded and a number of diverse and conflicting demands have arisen which require controlled allocation of critical facility resources (file space and central processor time). We have implemented user-oriented policies in trying to give users the greatest latitude possible to pursue their research consistent with fairly meeting our responsibilities in managing SUMEX as a national resource.

We have described the details of our allocation procedures in earlier reports. These have been implemented to attempt to maintain the 40:40:20 balance in system use between Stanford, National, and staff communities. The initial complement of user projects justifying the SUMEX resource was centered to a large extent at Stanford. As the number of national has grown, so has the Stanford group of projects matured and in practice the 40:40 split between Stanford and non-Stanford projects is not ideally

realized (see Appendix B). Our job scheduling controls bias the allocation of CPU time based on percent time consumed relative to the time allocated over the 40:40:20 community split. The controls are "soft" however in that they do not waste computer cycles if users below their allocated percentages are not on the system to consume the cycles. The operating disparity in CPU use to date reflects a substantial difference in demand between the Stanford community and the developing national projects, rather than inequity of access. For example, the Stanford utilization is spread over a large part of the 24-hour cycle, while national-AIM users tend to be more sensitive to local prime-time constraints. (The 3-hour time zone phase shift across the continent is of substantial help in load balancing.) During peak times under the new overload controls, the Stanford community still experiences mutual contentions and delays while the AIM group has relatively open access to the system. For the present, we propose to continue our policy of "soft" allocation enforcement for the fair split of resource capacity.

Our system also categorizes users in terms of access privileges. These comprise fully authorized users, pilot projects, guests, and network visitors in descending order of system capabilities. We want to encourage bona fide medical and health research people to experiment with the various programs available with a minimum of red tape while not allowing unauthenticated users to bypass the advisory group screening procedures by coming on as guests. So far we have had relatively little abuse compared to what other network sites have experienced, perhaps on account of the personal attention that senior staff gives to the logon records, and to other security measures. However, the experience of most other computer managers behooves us to be cautious about being as wide open as might be preferred for informal service to pilot efforts and demonstrations. We will continue developing this mechanism in conjunction with management committee policy decisions.

We have actively encouraged mature projects to apply for their own machine resources in order to preserve the SUMEX-AIM resource for new AI applications. In the recent past, several projects have submitted proposals for such facilities including DENDRAL (see Section 9.1.3 on page 149). In spite of favorable reviews of the research project itself (resulting in a 3-year renewal), the study section did not want to see the DENDRAL project divert its energies to run a separate machine resource. Rather they felt such an augmentation should be coordinated and implemented by the SUMEX resource in conjunction with the DENDRAL group. Such a relationship is feasible in the case of the local DENDRAL project and we feel can serve as a model for further distribution of resources to advanced projects. We cannot effectively operate such resources for all the projects in our community but through experimentation with new machines, we can lay the groundwork for packaged systems that other groups may be able to acquire and easily operate. This mandate through the DENDRAL review is one of the bases for our long term plans for the coming renewal period.

Appendix FLISP Address Space Limitations

In recent years, the program address space limitations imposed by the architecture of the PDP-10/20 systems have been increasingly felt in building large knowledge-based systems for biomedicine and in other application areas. Each user has access to a 256K 36-bit virtual address space (slightly more than 1M byte). For many conventional programs, this is adequate but the large language and program structures required for expert systems easily consume this space.

Current systems have used many approaches to compress their address space requirements including compiling established static code so it can be swapped between the main LISP space and an inferior fork and reorganizing dynamic code and data structures so they can be swapped between memory and hash-coded files. For example, space is now a critical problem for GUIDON because it is itself a large system built on top of another large system, MYCIN. In MYCIN, the dictionary, tables of facts (drugs/organism relations), and static properties that consume string space have already been moved off to disk in the form of hash files. In GUIDON, even this is not enough; MYCIN's rules must be hashed as well. For the short term, it appears that more of GUIDON's code will have to be non-resident ("recognized files"), thus trading time for space. Since response time is crucial for consultative programs, this trade-off is not acceptable.

Early in the development of Internist-I it became obvious that the 18 bit address space of INTERLISP imposed a severe limitation on the size of the knowledge base. The limit was on both atom and list space. To make matters worse there was no room left for the dynamic data structures (mostly lists) that are established by the diagnostic program. To get around this problem the INTERNIST group invested approximately 2 man years to develop a disk-oriented knowledge base that fetched and overlayed knowledge structures on demand. As a result all but the most trivial changes on knowledge structures are prohibitive, the system is not portable, and they still see an occasional case for which there is insufficient list space to be used by the diagnostic program.

Similar problems are anticipated in the development of Internist-II. The plan, at present, is to employ LISP hash files for the larger and/or infrequently accessed structures.

In both AGE and Meta-DENDRAL, it is not possible to load all the information on the system files into a single save file. This is handled by having different specialized environments that contain different system information, e.g., system execution and system development. In Meta-DENDRAL, all of the executing code will not fit in a single address space, so a system of selective loading is used based on dynamic demand. This reduces memory requirements for code but increases system overhead. In addition, DENDRAL has used a greatly stripped down version of LISP (also

used by INTERNIST) in order to have sufficient data space to handle meaningful problems. They are still constrained in problem complexity by the limited space to store data structures.

Similarly in MOLGEN, the address space in INTERLISP was sufficiently tight that the knowledge base would not fit in core, even at a very early stage in the project. To remedy this, they added a "virtual memory" system to the Units representation system which paged from a disk file on a demand basis. This patch basically made the PDP10 usable at a cost in execution time.

While the 18-bit address limit has not stopped research, it has stifled it by increasing overhead and causing users to scale down the scope of their research efforts. In order to minimize the cost of knowledge-base and program overlays, each project has had to tune their approach to the particular program structure. Even fairly modest ambitions push tolerance and system capacity to the limits. Much effort has gone into solving this problem in the ARPANET INTERLISP community. Address extensions for the PDP-10/20 class machines (including Foonly, Inc. machines) based on memory segmentation schemes do not lend themselves to a LISP environment since there is no intrinsic difference between program and data and the added overhead of keeping track of the extended address constructs with software becomes prohibitive. Thus, the solutions under active consideration include moving either to general purpose machines with larger logical address spaces (e.g., Prime or DEC VAX) or to special purpose LISP machines.

One of our objectives for the renewal period is to add facilities to the SUMEX-AIM resource that will provide a uniform and effective solution to these problems.

Appendix GAI Handbook Outline

E. A. Feigenbaum and A. Barr
Computer Science Department
Stanford University

This is a list of the Chapters in the Handbook. Articles in the first eight Chapters are expected to appear in Volume I. A tentative list of the all of articles in each Chapter follows.

- I. Introduction
- II. Search
- III. Representation of Knowledge
- IV. Natural Language Understanding
- V. Speech Understanding
- VI. AI Programming Languages
- VII. Applications-oriented AI Research: Science
- VIII. Applications-oriented AI Research: Medicine
- IX. Applications-oriented AI Research: Education
- X. Automatic Programming
- XI. Information Processing Psychology
- XII. Theorem Proving
- XIII. Vision
- XIV. Robotics
- XV. Learning and Inductive Inference
- XVI. Planning, Reasoning, and Problem Solving

I. INTRODUCTION

- A. The AI Handbook (intent, audience, style, use, outline)
- B. Overview of AI
- C. History of AI
- D. An Introduction to the AI Literature

II. Search

- A. Overview
- B. Problem representation
 - 1. State-space representation
 - 2. Problem-reduction representation
 - 3. Game trees
- C. Search methods
 - 1. Blind state-space search
 - 2. Blind AND/OR graph search
 - 3. Heuristic state-space search
 - a. Basic concepts in heuristic search
 - b. A*: optimal search for an optimal solution
 - c. Relaxing the optimality requirement
 - d. Bidirectional search
 - 4. Heuristic search of an AND/OR graph
 - 5. Game tree search
 - a. Minimax
 - b. Alpha-beta pruning
 - c. Heuristics in game tree search
- D. Example search programs
 - 1. Logic Theorist
 - 2. GPS
 - 3. Gelernter's geometry theorem-proving machine
 - 4. Symbolic integration programs
 - 5. STRIPS
 - 6. ABSTRIPS

III. Representation of Knowledge

- A. Issues and problems in representation theory
- B. Survey of representation techniques
- C. Representation schemes
 - 1. Logic
 - 2. Procedural representations
 - 3. Semantic networks
 - 4. Production systems
 - 5. Direct (analogical) representations
 - 6. Semantic primitives
 - 7. Frames and scripts

IV. Natural Language Understanding

- A. Overview - History and issues
- B. Early attempts at mechanical translation
- C. Grammars
 - 1. Review of formal grammars
 - 2. Transformational grammars
 - 3. Systemic grammars
 - 4. Case grammars
- D. Parsing
 - 1. Overview of parsing techniques
 - 2. Augmented transition nets, Woods
 - 3. CHARTS - The GSP system
- E. Text generating systems
- F. Natural language processing systems
 - 1. Early NL systems
 - 2. Wilks' machine translation work
 - 3. MARGIE
 - 4. LUNAR
 - 5. SHRDLU
 - 6. SAM and PAM
 - 7. LIFER

V. Speech Understanding Systems

- A. Overview
- B. Some early ARPA speech systems
 - 1. DRAGON
 - 2. HEARSAY I
 - 3. SPEECHLIS
- C. Recent Speech Systems
 - 1. HARPY
 - 2. HEARSAY II
 - 3. HWIM
 - 4. SRI-SDC System

VI. AI Programming Languages

- A. Historical overview
- B. AI programming language features
 - 1. Overview and comparison
 - 2. Data structures
 - 3. Control structures
 - 4. Pattern matching
 - 5. Programming environment
- C. Major AI programming languages
 - 1. LISP
 - 2. PLANNER and CONNIVER
 - 3. QLISP
 - 4. SAIL
 - 5. POP-2

VII. Applications-oriented AI Research: Science and Mathematics

- A. Overview
- B. TEIRESIAS - Issues in expert systems design
- C. Applications in chemistry
 - 1. Applications in chemical analysis
 - 2. The DENDRAL Programs
 - a. DENDRAL
 - b. CONGEN and its extensions
 - c. Meta-DENDRAL
 - 4. CRYNALIS
 - 5. Applications in organic synthesis
- D. Applications in mathematics
 - 1. MACSYMA
 - 2. AM
- F. Miscellaneous science applications research
 - 1. The SRI Computer-Based Consultant
 - 2. PROSPECTOR

VIII. Applications-oriented AI Research: Medicine

- A. Overview
- B. Medical systems
 - 1. MYCIN
 - 2. CASNET
 - 3. INTERNIST
 - 4. Present Illness Program
 - 5. Digitalis Advisor
 - 6. IRIS

IX. Applications-oriented AI Research: Education

- A. Historical overview
- B. Issues in ICAI systems design
- C. ICAI Systems
 - 1. SCHOLAR
 - 2. WHY
 - 3. SOPHIE
 - 4. WEST
 - 5. WUMPUS
 - 6. BUGGY
 - 7. EXCHECK

X. Automatic Programming

- A. Overview - Methods of program specification
- B. Basic approaches
- C. AP Systems
 - 1. PSI
 - 2. SAFE
 - 3. Programmer's Apprentice
 - 4. PECOS
 - 5. DAEDALUS
 - 6. PROTOSYSTEM-1
 - 7. NLPQ
 - 8. LIBRA - Program Optimization

XI. Information Processing Psychology

- A. Overview
- B. GPS
- C. Cognitive development
- D. EPAM
- E. Semantic network models
 - a. Quillian's network
 - b. LNR's MEMOD
 - c. HAM
 - d. ACT
- F. Belief systems

XII. THEOREM PROVING

- A. Overview
- B. Logic
- C. Resolution theorem proving
 - 1. Basic resolution method
 - 2. Syntactic ordering strategies
 - 3. Semantic and syntactic refinement
- D. Non-resolution theorem proving
 - 1. Overview
 - 2. Natural deduction
 - 3. Boyer-Moore
 - 4. LCF
- E. Applications of theorem proving
 - 1. Use in question answering
 - 2. Use in problem solving
 - 3. Theorem proving programming languages
 - 4. Man-machine theorem proving
 - 5. Use in automatic programming
- F. Proof checkers

XIII. VISION

- A. Overview
- B. Image-level processing
 - 1. Overview
 - 2. Edge detection
 - 3. Texture
 - 4. Region growing
 - 5. Overview of pattern recognition
- C. Spatial-level processing
 - 1. Overview
 - 2. Stereo information
 - 3. Shading
 - 4. Motion
- D. Object-level processing
 - 1. Overview
 - 2. Generalized cones and cylinders
- E. Scene level processing

- F. Vision systems
 - 1. Polyhedral or Blocks World vision
 - a. Overview
 - b. COPYDEMO
 - b. Guzman
 - c. Falk
 - d. Waltz
 - e. Navatya
 - 2. Robot vision systems
 - 3. Perceptrons
- XIV. Robotics
 - A. Overview
 - B. Robot planning and problem solving
 - C. Arms
 - D. Present-day industrial robots
 - E. Robotics programming languages
- XIII. Learning and Inductive Inference
 - A. Overview
 - B. Simple inductive tasks
 - 1. Sequence extrapolation
 - 2. Grammatical inference
 - C. Pattern recognition
 - 1. Character recognition
 - 2. Other recognition tasks
 - D. Learning rules and strategies of games
 - 1. Formal analysis
 - 2. Examples of game-learning programs
 - E. Single concept formation
 - F. Multiple concept formation: Structuring a domain (AM, Meta-DENDRAL)
 - G. Interactive cumulation of knowledge (TEIRESIAS)
- XIV. Problem Solving, Planning & Reasoning by Analogy
 - A. Overview of problem solving
 - B. Planning
 - 1. Overview
 - 2. STRIPS (see IID5)
 - 3. ABSTRIPS (see IID6)
 - 4. NOAH
 - 5. HACKER
 - 6. INTERPLAN
 - 7. Rieger's causal reasoning system
 - 8. Rutgers work
 - 7. QA3 (see IXE1)
 - C. Reasoning by analogy
 - 1. Overview
 - 2. Evans's ANALOGY program
 - 3. ZORBA
 - 4. Winston's learning system
 - D. Constraint relaxation
 - 1. Waltz
 - 2. REF-ARF
 - E. Game playing

Appendix HMAINSAIL System Demonstration

As of July 30, 1979, the MAINSAIL project has successfully designed, demonstrated, and documented an ALGOL-like language system for machine-independent software design. This system includes the compiler, code generators, and run-time support for a range of target machine environments including TENEX, TOPS-20, TOPS-10, RT-11, and RSX-11. The designs for other environments have been studied but resources have not allowed more extensive implementations. Within Council-approved funding and manpower limits and the AI charter of the SUMEX resource, we do not have access to the more extensive resources that would be required to continue effective development and export of this system beyond this initial research and demonstration phase. The principal individuals involved (Messrs. Wilcox and Jirak and Ms. Dageforde) have formed a small private company, XIDAK, to support and continue development of MAINSAIL under license from Stanford University. XIDAK has almost completed a VAX implementation of MAINSAIL and is pursuing interests from a growing group of potential users, including a microprogrammed implementation for the PERQ computer. The following is a brief summary of recent work in this final demonstration phase of the MAINSAIL effort. Detailed reports on the language manual and design description can be found in references 14 and 15.

- 1) The compiler has undergone major reexamination and improvement with a substantial reduction in the size of data structures. As a result, it is now able to run on 16-bit machines with small address spaces (e.g., 32K words).
- 2) The runtime systems were thoroughly reexamined for optimizing execution efficiency and memory utilization. The garbage collection facility, used in the dynamic storage allocation system, was also substantially improved.
- 3) A new approach to code generation was introduced utilizing tree structures for the intermediate representation, rather than the more primitive triples or quadruples.
- 4) Facilities for managing "module libraries" of executable MAINSAIL modules were implemented.
- 5) At the conclusion of the demonstration phase, there were three sites using the TENEX version, six using the TOPS-10 version, and five using the TOPS-20 version.
- 6) A research project based on MAINSAIL is underway, aimed at an efficient program execution and development environment implemented on a microcoded "MAINSAIL machine" which directly executes a tailor-made MAINSAIL instruction set. This is the basis of Wilcox's Ph.D. thesis.

Appendix I

AIM Management Committee Membership

The following are the membership lists of the various SUMEX-AIM management committees at the present time:

AIM Executive Committee:

LEDERBERG, Joshua, Ph.D. President The Rockefeller University 1230 York Avenue New York, New York 10021 (212) 360-1234, 360-1235	(Chairman)
AMAREL, Saul, Ph.D. Department of Computer Science Rutgers University New Brunswick, New Jersey 08903 (201) 932-3546	
BAKER, William R., Jr., Ph.D. Biotechnology Resources Program National Institutes of Health Building 31, Room 5B43 9000 Rockville Pike Bethesda, Maryland 20205 (301) 496-5411	(Exec. Secretary)
FEIGENBAUM, Edward, Ph.D. Principal Investigator - SUMEX Department of Computer Science Margaret Jacks Hall, Room 216 Stanford University Stanford, California 94305 (415) 497-4079	
LINDBERG, Donald, M.D. 605 Lewis Hall University of Missouri Columbia, Missouri 65201 (314) 882-6966	(Adv Grp Member)
MYERS, Jack D., M.D. School of Medicine Scaife Hall, 1291 University of Pittsburgh Pittsburgh, Pennsylvania 15261	

SHORTLIFFE, Edward H., M.D., Ph.D.
Co-Principal Investigator - SUMEX
Division of General Internal Medicine, TC117
Stanford University Medical Center
Stanford, California 94305
(415) 497-5821

AIM Advisory Group:

LINDBERG, Donald, M.D.	(Chairman)
605 Lewis Hall	
University of Missouri	
Columbia, Missouri 65201	
(314) 882-6966	
AMAREL, Saul, Ph.D.	
Department of Computer Science	
Rutgers University	
New Brunswick, New Jersey 08903	
(201) 932-3546	
BAKER, William R., Jr., Ph.D.	(Exec. Secretary)
Biotechnology Resources Program	
National Institutes of Health	
Building 31, Room 5B43	
9000 Rockville Pike	
Bethesda, Maryland 20205	
(301) 496-5411	
FEIGENBAUM, Edward, Ph.D.	(Ex-officio)
Principal Investigator - SUMEX	
Department of Computer Science	
Margaret Jacks Hall, Room 216	
Stanford University	
Stanford, California 94305	
(415) 497-4079	
LEDERBERG, Joshua, Ph.D.	
President	
The Rockefeller University	
1230 York Avenue	
New York, New York 10021	
(212) 360-1234, 360-1235	
MINSKY, Marvin, Ph.D.	
Artificial Intelligence Laboratory	
Massachusetts Institute of Technology	
545 Technology Square	
Cambridge, Massachusetts 02139	
(617) 253-5864	
MOHLER, William C., M.D.	
Associate Director	
Division of Computer Research and Technology	
National Institutes of Health	
Building 12A, Room 3033	
9000 Rockville Pike	
Bethesda, Maryland 20205	
(301) 496-1168	

- MYERS, Jack D., M.D.
 School of Medicine
 Scaife Hall, 1291
 University of Pittsburgh
 Pittsburgh, Pennsylvania 15261
 (412) 624-2649
- PAUKER, Stephen G., M.D.
 Department of Medicine - Cardiology
 Tufts New England Medical Center Hospital
 171 Harrison Avenue
 Boston, Massachusetts 02111
 (617) 956-5910
- SHORTLIFFE, Edward H., M.D., Ph.D. (Ex-officio)
 Co-Principal Investigator - SUMEX
 Division of General Internal Medicine, TC117
 Stanford University Medical Center
 Stanford, California 94305
 (415) 497-5821
- SIMON, Herbert A., Ph.D.
 Department of Psychology
 Baker Hall, 339
 Carnegie-Mellon University
 Schenley Park
 Pittsburgh, Pennsylvania 15213
 (412) 578-2787 or 578-2000

Stanford Community Advisory Committee:

FEIGENBAUM, Edward, Ph.D. (Chairman)

Department of Computer Science
Margaret Jacks Hall, Room 216
Stanford University
Stanford, California 94305
(415) 497-4079

SHORTLIFFE, Edward H., M.D., Ph.D.

Co-Principal Investigator - SUMEX
Division of General Internal Medicine, TC117
Stanford University Medical Center
Stanford, California 94305
(415) 497-5821

DJERASSI, Carl, Ph.D.

Department of Chemistry, Stauffer I-106
Stanford University
Stanford, California 94305
(415) 497-2783

LEVINTHAL, Elliott C., Ph.D.

Department of Genetics, S047
Stanford University Medical Center
Stanford, California 94305
(415) 497-5813